Colored minority games

Matteo Marsili[(1)] and Maurizio Piai[(2)]

[(1)] *Istituto Nazionale per la Fisica della Materia (INFM), Trieste-SISSA Unit, V. Beirut 2-4, Trieste I-34014*
[(2)] *SISSA-ISAS, V. Beirut 2-4, Trieste I-3401 and Istituto Nazionale di Fisica Nucleare (INFN),sez. Trieste*

1 Introduction

The Minority Game (MG) [1] provides an ideal playground for studying interacting agent systems such as financial markets [2]. Its stationary states have indeed been understood in great detail [3–5] whose conclusions have been recently confirmed by a fully dynamic theoretical approach [6]. Such a state of the art allows one to address interesting issues on the way in which such complex systems behave.

The MG is a largely oversimplified model of a systems of adaptive agents. This characteristic is what makes the MG analytically tractable thus allowing one to derive a coherent picture of its collective behavior. Then, it is important to understand how this picture changes as one moves towards more realistic cases. Several investigations have indeed aimed at relaxing the many unrealistic features of the MG: The effect of non-adaptive agents (also called producers or hedgers) has been studied in Ref. [2]; the role of market impact was addressed in Refs. [3,7,8] and differently weighted agents were discussed in Ref. [9]. Agents who can refrain from playing, if that is not convenient have also been considered in Ref. [10–12]. The analytic theory offers insights on these issues which allows one to form a far more complete picture than that resulting only from numerical simulations [13].

An unrealistic feature of the MG is that agents trade on the same time-scales and that different events also occur with the same frequency. Our aim in this note is to study how the collective behavior changes if one goes beyond such simplifying assumptions.

The MG is based on a *syncronous* dynamics. Every agent, at each time trades in the market. In the real world agents trade on different frequencies; some everyday, even several times in a day, and some once a week, a month or an year. The distribution of trading time-scales across agents may likely be spread over several decades, which may offer an explaination for the presence of long range correlations in the volume of financial activity [11]. How does the behavior of MG changes if we account for strategies acting on different, even widely spread, frequencies?

We find that the qualitative behavior of the system is not affected. Still a phase transition between an informationally efficient phase and an inefficient phase exist. We quantify how the critical point where the phase transition occurs depends on the distribution.

The second question we ask concerns the frequency with which events occur. In the world of the MG one of many possible events occur in each period. Agents observe these events and indeed they adopt strategies which provide actions conditional to events. All events in the MG occur with the same probability. In section 4 we study a situation where events occur with different probabilities. Our discussion builds up on results derived in [14]. In particular we address the issue of information efficiency.

When the market is not informationally efficient, one may investigate market predictability as a function of the frequency of events. We find that markets are more predictable when rare events occur than when typical events occur. In particular the asymmetry of market outcome in a given state is inversely proportional to the frequency of that state. This gives a frequency-dependent characterization of market (in)efficiency.

When comparing these results with real market data, one faces

the problem of identifying relevant events. We discuss a simple example based on FX data. Average returns of the FX rate are conditioned to events which, as in the the original MG, encode the information on the M most recent market movements. Our evidence is not strong but it points in the right direction: The market is more predictable when rare events occur.

2 The Minority Game

The MG depicts a market with N adaptive agents or speculators. The market can be in one of P states, labelled by μ in the following. Each trader i has two personal trading strategies, labeled by a spin variable s_i, which prescribe an action $a^{\mu}_{s_i,i}$ for each state μ. These trading strategies are drawn at random from a certain distribution (see later) and assigned to agents.

The game is repeated many times; at time t the state $\mu(t)$ is drawn from a distribution ρ^{μ} at each time and agents try to estimate, on the basis of past observations, which of their strategies is the best one. To do this, agents assign a *score* $U_{s,i}(t)$ to each strategy $s = \pm 1$. The larger the score of a strategy, the more likely it will be played by the agent: If $s_i(t)$ is the strategy played by agent i at time t we assume that

$$\mathrm{Prob}[s_i(t) = s] \propto \exp\left[\Gamma U_{s,i}(t)\right] \tag{1}$$

where Γ is a *learning rate*. Each agent monitors the scores $U_{s,i}(t)$ of each of her strategies s by

$$U_{s,i}(t+1) = U_{s,i}(t) - a^{\mu(t)}_{s,i} A(t)/N, \quad A(t) = \sum_j a^{\mu(t)}_{s_j(t),j}. \tag{2}$$

This dynamics implies that agents prefer to take actions with the opposite sign to that of $A(t)$. Indeed they reward by increasing $U_{s,i}$ the score of a strategy s for which $a^{\mu(t)}_{s,i} A(t) < 0$. If only binary actions $a^{\mu}_{s,i} = \pm 1$ are allowed, then $A(t)$ has the sign of

the action taken by the majority. Hence the minority "wins". We refer the interested reader to Refs. [1,9,15] for deeper discussions on the interpretation of this dynamics.

The relevant parameter [16] is the ratio

$$\alpha = P/N$$

between the "information complexity" P and the number of agents, and the key quantity we shall look at is

$$\sigma^2 = \langle A^2 \rangle$$

where $\langle \ldots \rangle$ is a time average in the stationary state. σ^2 is a measure of the inefficiency of agents' coordination because. Hence optimal states correspond to minima of σ^2.

A further interesting quantity is the predictability

$$H = \sum_{\mu=1}^{P} \rho^\mu \langle A|\mu \rangle^2 \tag{3}$$

of the market. Here $\langle \ldots |\mu \rangle$ is a conditional average in the state μ. If the average $\langle A|\mu \rangle$ of $A(t)$ conditional to $\mu(t) = \mu$ is non-zero, then the sign of $A(t)$ is statistically predictable if μ is known. This is why H measures predictability.

Previous works have mainly focused on strategies drawn from the distribution

$$P(a^\mu_{s,i}) = \frac{1}{2}\delta(a^\mu_{s,i} + 1) + \frac{1}{2}\delta(a^\mu_{s,i} - 1) \tag{4}$$

(a correlation between the two trading strategies has also been considered[2]) and on the event distribution

$$\rho^\mu = \frac{1}{P}. \tag{5}$$

The collective behavior of the market is the following: For $\alpha > \alpha_c$ the market is predictable ($H > 0$), the stationary state is unique and independent of the learning rate Γ. As α decreases from large values, also σ^2 decreases signalling that agents manage to coordinate more and more efficiently. At $\alpha = \alpha_c$ a phase transition occurs. Indeed for $\alpha < \alpha_c$ the market is unpredictable ($H = 0$), the stationary state is not unique but rather depends both on the initial conditions and on the learning rate. For symmetric initial conditions ($U_{s,i}(0) = 0\ \forall s, i$), σ^2 reaches a minimum at α_c and then it starts raising as α decreases below α_c. In this region σ^2 is also an increasing function of Γ. The inefficiency σ^2 also decreases as a function of the asymmetry of initial conditions. For strongly asymmetric initial conditions (e.g. $U_{+,i}(0) \gg U_{-,i}(0)$, $\forall i$) σ^2 is greatly reduced and it decreases with α.

Such a complex behavior can be captured and described quantitatively by a statistical mechanics approach [5] based on two steps: First one studies the stationary state properties of the system and secondly one studies the stochastic dynamics. Here we discuss how this picture changes when either of Eqs. (4,5) is modified to allow for agents acting or events occurring on different frequencies.

3 Agents trading on different time-scales

Let us assume that agent i plays with a frequency f_i. More precisely we assume that

$$P(a^\mu_{s,i}) = f_i \left[\frac{1}{2}\delta(a^\mu_{s,i}+1) + \frac{1}{2}\delta(a^\mu_{s,i}-1)\right] + (1-f_i)\delta(a^\mu_{s,i}). \quad (6)$$

This means that on a fraction $1-f_i$ of the events, agent i does not trade ($a^\mu_{s,i} = 0$). Different agents play with different frequencies, and are not supposed to respect any rigid time-table. We assume f_i to be fixed at the beginning of the game, sorting it randomly on the basis of a given distribution $\nu(f)$. In order to describe the

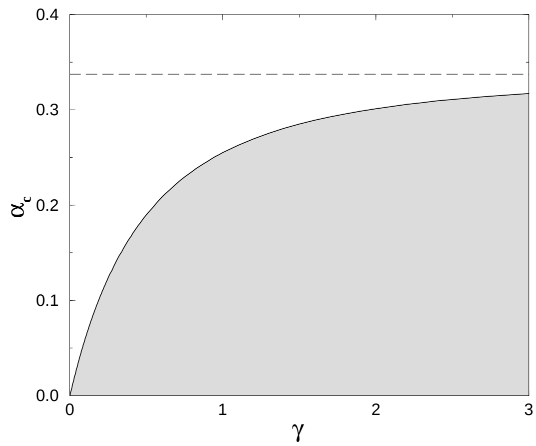

Fig. 1. Phase diagram in the (α, γ) plane. The shaded region below the curve $\alpha_c(\gamma)$ is the efficient (symmetric) phase. The dashed line corresponds to $\alpha_c(\infty) \approx 0.3374\ldots$

lack of a single characteristic time-scale, we will mainly consider the distribution

$$\nu(f) = \gamma\, f^{\gamma-1}, \tag{7}$$

so that the system is now described by two parameters γ and α. In the limit $\gamma \to \infty$ we recover the case in which all players are supposed to act at each run of the game. Following Ref. [5], we find that the stationary state "magnetizations" $m_i = \langle s_i \rangle$ are coincide with the minima of H where, in Eq. 3,

$$\langle A|\mu\rangle = \sum_{i=1}^{N} \left[\frac{a_{+,i}^{\mu} + a_{-,i}^{\mu}}{2} + m_i \frac{a_{+,i}^{\mu} - a_{-,i}^{\mu}}{2} \right].$$

The properties of the stationary states are then accessible using standard methods of statistical mechanics of disordered systems. The approach follows exactly the same steps as those of Refs. [3,2]. As detailed in the appendix, the main difference concerns the disorder average of the replicated partition function that is now taken with Eq. (6) and not with Eq. (4). Furthermore one has to average over the distribution $\nu(f)$. Details of the calculations are reported in the appendix.

The system behaves in a non-trivial way, presenting a phase transition in the (γ, α) plane as shown in figure 1). This separates an

informationally inefficient phase with $H > 0$ – for $\alpha > \alpha_c(\gamma)$ – from an efficient phase ($H = 0$) for $\alpha < \alpha_c(\gamma)$. As expected when $\gamma \to \infty$ we recover the value of the critical threshold $\alpha_c \approx 0.3374\ldots$ of the mono-chromatic MG [3]. For finite values of α we find a smaller value of α_c. This is consistent with the observation that a smaller number of agents are active in the market at each time.

When $\gamma \to 0$ the critical threshold $\alpha_c(\gamma)$ vanishes. This is consistent with the observation that in such a market the effective number of trader is a vanishing fraction of N.

Concerning agent's behavior we find that the distribution of magnetization m conditional to the frequency f is given by:

$$P(m|f) = \frac{1}{2}\phi(f)\left[\delta(m+1) + \delta(m-1)\right] + \frac{\zeta f}{\sqrt{2\pi}} e^{-(\zeta f m)^2/2} \quad (8)$$

where $\zeta = \zeta(\alpha, \gamma)$ is a parameter which depends on α and γ (see appendix) and

$$\phi(f) = \mathrm{erfc}\left(\frac{\zeta f}{\sqrt{2}}\right).$$

is the probability that an agent with $f_i = f$ is *frozen*, i.e. that he/she plays always the same strategy ($m_i = \pm 1$). Hence agents who play less frequently (small f_i) are more likely to be frozen. This is reasonable because it is the market impact [4] which makes agents switch between strategies and agents who trade less frequently are less affected by the market impact.

Most of the *noise traders* – i.e. of the agents with $m_i \approx 0$ – are frequent traders: It is not difficult to see that the probability that a trader with $m_i \approx 0$ has frequency $f_i = f$ is $\mathrm{Prob}(f_i = f | m_i \approx 0) \propto f^\gamma$.

The main conclusion of this section is that, even with a broad distribution of trading time-scales across agents, the basic picture

of the behavior of the MG remains that of the mono-chromatic MG. We also see that frequent traders are responsible for market volatility.

4 Broadly distributed frequency of events

In this section we address the behavior of a market, as described by the MG, when the basic events have a broad distribution of frequencies. In other words, instead of Eq. (5) we consider

$$\rho^\mu = \frac{1}{P}\tau^\mu \tag{9}$$

with τ^μ distributed with a pdf $\theta(\tau)$ such that

$$\int_0^\infty \theta(\tau)\tau d\tau = 1$$

which ensures normalization of ρ^μ in Eq. (9). Now there will be events which occur more frequently than others with a spread of frequencies which depends on the distribution $\theta(\tau)$.

The solution of the MG with a generic ρ^μ has been discussed in detail in Ref. [14]. Here we pick the main results:

The predictability can be written as

$$H = \frac{1}{P}\sum_{\mu=1}^{P} \tau^\mu \langle A|\mu\rangle^2 \cong \int_0^\infty d\tau\theta(\tau)\tau\langle A|\tau\rangle^2$$

where $\langle A|\tau\rangle^2$ is the average predictability of events with frequency τ. This quantity can be read off from Eq. (C2) of Ref. [14], which reads:

$$\frac{H}{N} = \frac{1+Q}{2}\int_0^\infty d\tau\theta(\tau)\frac{\tau}{(1+\tau\chi)^2}$$

where χ and Q are given in appendix 4.1. From these equations we find that the typical asymmetry $|\langle A|\tau\rangle|$ that one can expect on events with frequency τ is

$$|\langle A|\tau\rangle| = \sqrt{\frac{1+Q}{2}\frac{1}{1+\tau\chi}}. \tag{10}$$

The quantity $|\langle A|\tau\rangle|$ measures the maximal excess return that a trader can possibly exploit by trading on events at time-scale τ. In a market where agents behave randomly ($m_i = 0$ for all i) we expect $|\langle A|\tau\rangle|$ to be independent of τ. Eq. (10) shows that traders activity exploits more heavily arbitrages on frequent events thus reducing $|\langle A|\tau\rangle|$. The result of this is that, for frequent events ($\tau\chi \gg 1$), the excess return is inversely proportional the probability τ/P of events.

Because the number of traders is finite and their ability of detecting and exploiting arbitrages is bounded, rare events ($\tau\chi \approx 1$) carry arbitrages which are left unexploited. There is a characteristic frequency $1/\chi$ below which the law $|\langle A|\tau\rangle| \sim 1/\tau$ levels off: For rare events $\rho^\mu < 1/(\chi P)$ arbitrage size remains constant, suggesting that traders are unable to optimize their behavior at a time resolution smaller than $1/\chi$.

In the limit $\alpha \to \alpha_c^+$ the market becomes efficient so $|\langle A|\tau\rangle| \to 0$. At the same time, however $\chi \to \infty$ which means that the $|\langle A|\tau\rangle| \sim 1/\tau$ behavior extends to extremely rare events.

Fig. 2 shows that numerical simulations are in full agreement with Eq. (10). This plot in particular allows one to measure χ from market data. This in turn provides a measure of the inefficiency of the market.

4.1 Efficiency in real markets

Let us discuss a practical application to the above results to real market data. There are two main sources of problems: First it

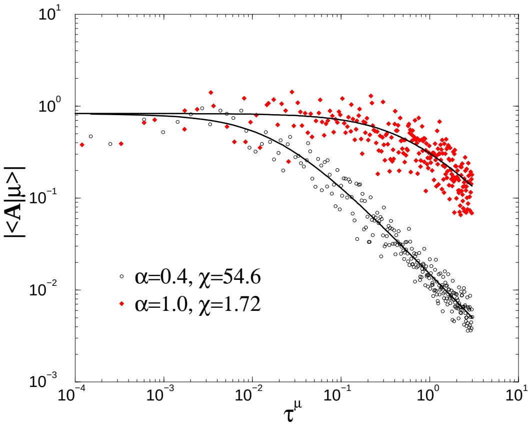

Fig. 2. Excess return $|\langle A|\mu\rangle|$ as a function of τ^μ from numerical simulations of sistems with $P = 256$ events, $\alpha = 0.4$ and 1.0. The simulation used $\theta(\tau) = a\tau^{-1/3}$ for $\tau < \tau_>$. The full lines are numerical fits with the functional form $A/(1 + \chi\tau)$.

is not clear how to identify a set of exclusive events. Events are specified in the very definition of the MG but it is not clear what are their counterparts in a real market.

Second one faces the problem of handling a finite data set. Suppose that we have T data points and that, having identified a set of events, we define $\mu(t)$ as the label of the event which occurs in observation $t = 1, \ldots, T$. Then our estimate $\widehat{\langle A|\mu\rangle}$ of the market return conditional to event μ will be

$$\widehat{\langle A|\mu\rangle} = \frac{1}{T^\mu}\sum_{t=1}^{T}\delta_{\mu,\mu(t)}A(t), \qquad T^\mu = \sum_{t=1}^{T}\delta_{\mu,\mu(t)}.$$

This will however be affected by statistical errors of the order

$$\delta\widehat{\langle A|\mu\rangle} \simeq \sqrt{V^\mu/T^\mu},$$

where V^μ is the sample variance, conditional to $\mu(t) = \mu$. One should then require that $\widehat{\langle A|\mu\rangle} \gg \delta\widehat{\langle A|\mu\rangle}$, i.e. that the results are statistically significant. This means that market data should be classified in a number of events as small as possible. However grouping different market conditions in the same state μ may average out relevant information.

We illustrate these points with a specific example. We analyzed

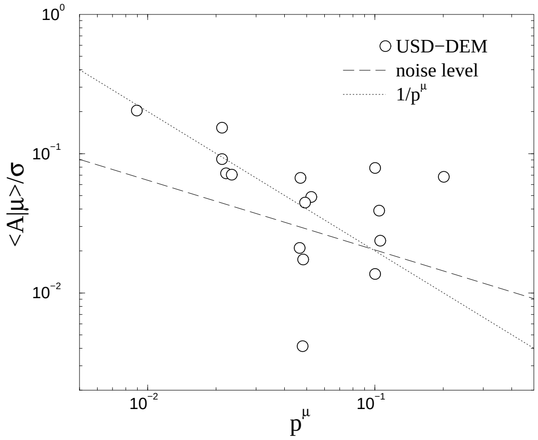

Fig. 3. Conditional excess returns $|\langle A|\mu\rangle|$ as a function of the frequancy f^μ in FX markets. For a definition of events E^μ see Refs. [17,18].

a data set of FX data on the USD-DEM exchange rate for the year 1993 [17]. Considering a typical time interval of 5 minutes, this gives $T \approx 10^5$ points. We define event μ, as in the original MG, as the encoding of the M most recent signs of the FX rate increments [18]. Taking $M = 4$, i.e. $P = 16$ events, we find the results of Fig. 3. The error on $\langle \hat{A|}\mu\rangle$ is reported as a dashed line in Fig. 3. The evidence for the law $|\langle A|\mu\rangle| \sim 1/\tau^\mu$ is very weak and hardly emerges from the noise level. In other words, one cannot rule out the hypothesis that $\langle A|\mu\rangle = 0$ and that what is seen is just fluctuations due to finite sample size. Furthermore there is no evidence of saturation for rare events. This is consistent with the fact that FX markets are very efficient and hence one expects a fairly large value of χ.

The present statistical resolution does not allow for sharp statements on the inefficiency of FX-markets. It serves however to illustrate the practical issues that such an empirical analysis raises.

Appendix

We give here some results of the algebraic manipulation of the model defined in section 3, and the final equation, to be solved numerically. In order to study the minima of H, as usual one

introduces a fictitious temperature β and builds the partition function $Z(\beta) = \mathrm{Tr}_m e^{-\beta H}$. In order to deal with quenched disorder (i.e. with the heterogeneity of agents' trading strategies) we resort to the replica trick [19]. This amounts in averaging the parition function $Z^n(\beta)$ of n copies of the system over the disorder and then taking the limit $n \to \infty$. This allows one to compute the disorder average of $\log Z(\beta)$ which is a self-averaging quantity. It is easy to see [3] that the resulting free-energy is correctly described by a simple replica symmetric (RS) ansatz, i.e.

$$\frac{1}{N}\langle \log Z(\beta)\rangle \simeq -\beta f$$

where

$$f = \frac{\alpha}{2\beta}\Big[\ln\big(1+\frac{\beta\gamma}{\alpha(\gamma+2)}(Q-q)\big) + \frac{\beta(1+q)}{\alpha(\gamma+2)/\gamma+\beta(Q-q)}\Big]$$
$$+\frac{\alpha}{2\beta}[QR-qr] - \frac{1}{\beta}\int_0^1 \mathrm{d}f\,\gamma\, f^{\gamma-1}\left\langle \ln Tr_m \exp\big(-\beta f V_z(m|f)\big)\right\rangle_z,$$

where

$$\begin{aligned} Q &= \sum_{\imath=1}^{N} f_i^2\, m_i^2 / \sum_{\imath=1}^{N} f_i^2 \\ q &= \sum_{\imath=1}^{N} f_i^2\, m_i m_i' / \sum_{\imath=1}^{N} f_i^2 \\ V_z(m|f) &= -\frac{\gamma+2}{\gamma}\frac{\alpha\beta f}{2}(R-r)m^2 - \sqrt{(\gamma+2)\alpha r/\gamma}\,zm. \end{aligned}$$

Here m_i and m_i' refer to two distinct replicas of the system. Finally R and r arise as Langrange multipliers. The analysis of the saddle point equations follows the usual steps (see e.g. the appendics of Refs. [4,2]). The limit $\beta \to 0$ has to be taken in the end. The solution can be expressed in parametric form in terms of a variable ζ:

$$Q = q$$

$$Q = 1 - (\gamma+2)\int_0^1 df\, f^{\gamma+1}\left[\sqrt{\frac{2}{\pi}}\frac{e^{-f^2\zeta^2/2}}{f\zeta} + \left(1-\frac{1}{f^2\zeta^2}\right)\mathrm{erf}\left(\frac{f\zeta}{\sqrt{2}}\right)\right]$$

and

$$1+Q = \frac{\gamma+2}{\gamma}\frac{\alpha}{\zeta^2}.$$

The critical threshold α_c is as usual [3,4,2,5] derived imposing that the fraction of non-frozen agents is qual to α:

$$\int_0^1 [1-\phi(f)]\gamma f^{\gamma-1} df = \int_0^1 \mathrm{erf}\left(\frac{f\zeta}{\sqrt{2}}\right)\gamma f^{\gamma-1} df = \alpha.$$

These equations can be easily solved numerically.

The analytic solution of the MG with generic distribution ρ^μ is described in Ref. [14] and follows the same lines as above. The parameters $\chi = \beta(Q-q)/\alpha$ and Q are given by the solution of the set of equations:

$$\begin{aligned} Q &= 1 - \sqrt{\frac{2}{\pi}}\frac{e^{-\zeta^2/2}}{\zeta} - \left(1-\frac{1}{\zeta^2}\right)\mathrm{erf}\left(\frac{\zeta}{\sqrt{2}}\right) \\ \alpha(1+Q)\langle\frac{\tau^2}{[1+\tau\chi]^2}\rangle_\tau &= \left[\frac{\mathrm{erf}(\zeta/\sqrt{2})}{\chi\zeta}\right]^2 \\ \langle\frac{\tau\chi}{1+\tau\chi}\rangle_\tau &= \frac{\mathrm{erf}(\zeta/\sqrt{2})}{\alpha} \end{aligned}$$

where $\langle\ldots\rangle_\tau = \int_0^\infty d\tau\theta(\tau)\ldots$. These equations are valid in the symmetric phase ($\alpha > \alpha_c$). χ diverges at the critial point, where $\alpha_c = \mathrm{erf}(\zeta/\sqrt{2})$.

References

[1] Challet D. and Zhang Y.-C., Physica A **246**, 407 (1997); e-print adap-org/9708006

[2] Challet D., M. Marsili and Y.-C. Zhang, Physica A **276**, 284 (2000) e-print cond-mat/9909265

[3] D. Challet, M. Marsili and R. Zecchina, Phys. Rev. Lett. **84**, 1824 (2000), e-print cond-mat/9904392

[4] M. Marsili, D. Challet and R. Zecchina, Physica A **280**, 522 (2000), e-print cond-mat/9908480

[5] M. Marsili and D. Challet, Phys. Rev. E **64**, 056138 (2001).

[6] J. A. F. Heimel, A. C. C. Coolen, Phys. Rev. E **63**, 056121 (2001) and e-print cond-mat/0107600

[7] M. Marsili and D. Challet, Adv Compl. Sys. **3**-I (2001), e-print cond-mat/0004376

[8] M. Marsili, P. Mulet, F. Ricci-Tersenghi and R. Zecchina, Phys. Rev. Lett. **87**, 208701 (2001).

[9] D. Challet, A. Chessa, M. Marsili, Y.-C. Zhang, Quant. Fin. **1**, 168 (2001); e-print cond-mat/0011042

[10] P. Jefferies P, M.L. Hart, P.M. Hui, N.F. Johnson, Eur. Phys. J. B **20**, 493(2001).

[11] J.-P. Bouchaud, I. Giardina, M. Mezard, Quantitative Finance **1**, 212 (2001); e-print cond-mat/0012156

[12] D. Challet, M. Marsili and Y.-C. Zhang, Physica A **294**, 514 (2001); e-print cond-mat/0101326

[13] Many other modifications of the basic MG have been considered in the literature using approaches based mainly on numerical simulations. See `http://www.unifr.ch/econophysics/minority/` for a commented collection of papers.

[14] D. Challet and M. Marsili, Phys Rev. E **62**, 1862 (2000).

[15] M. Marsili, Physica A **299** 93 (2001).

[16] Savit R., Manuca R., and Riolo R., Phys. Rev. Lett. **82(10)**, 2203 (1999), e-print adap-org/9712006.

[17] The data set is the same as that used in S. Galluccio, G. Caldarelli, M. Marsili and Y.-C. Zhang, Physica A, **245**, 423, (1997). The time of the data set is “rescaled” with a nonlinear transformation in order to avoid periodic patterns (daily or weekly). Accordingly, more or less points are taken in periods of high or low activity, respectively. On average we take one point every 5 minutes or 100 values for each trading day.

[18] If $x(t)$ is the FX rate at time t, we define $\chi(t) = 1$ if $x(t) > 1.0003x(t-1)$ and $\chi(t) = 0$ otherwise. The event $E^\mu(t) = \{Y(t) = \mu\}$ is defined in terms of the random variable

$$Y(t) = \sum_{t'=t-M+1}^{t} \chi(t') 2^{t-t'}.$$

Then we measure conditional averages $x^\mu = \langle x(t+1) - x(t) | Y(t) = \mu \rangle$ and $\sigma^\mu = \sqrt{\langle [x(t+1) - x(t)]^2 | Y(t) = \mu \rangle}$. Finally we compute $A^\mu = x^\mu / \sigma^\mu$ and the frequency $f^\mu = \langle \delta_{Y(t),\mu} \rangle$.

[19] M. Mezard, G. Parisi, M. A. Virasoro, *Spin glass theory and beyond* World Scientific (1987).